

03/16/00
U.S. PTO

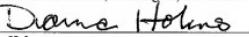
03-20-00

jc539 U.S. PTO
03/09/52085
03/16/00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

CERTIFICATE OF EXPRESS MAILING

I hereby certify that this paper and the documents and/or fees referred to as attached thereto are being deposited with the United States Postal Service on March 16, 2000 in an envelope as "Express Mail Post Office to Addressee" service under 37 CFR §1.10, Mailing Label Number EL54838485US, addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, DC 20231.


Dionna Holmes

Attorney Docket No.: CISCP141

First Named Inventor: Shmuel Shaffer

UTILITY PATENT APPLICATION TRANSMITTAL (37 CFR. § 1.53(b))

Assistant Commissioner for Patents
Box Patent Application
Washington, DC 20231

Duplicate for
fee processing

Sir: This is a request for filing a patent application under 37 CFR. § 1.53(b) in the name of inventors:
Shmuel Shaffer, James A. Aviani, Jr., Alexander G. Tweedy

For: METHODS AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC

Application Elements:

- 29 Pages of Specification, Claims and Abstract
 7 Sheets of **formal** Drawings
 9 Pages Combined Declaration and Power of Attorney

Accompanying Application Parts:

- Assignment and Assignment Recordation Cover Sheet (recording fee of \$40.00 enclosed)
 37 CFR 3.73(b) Statement by Assignee
 Information Disclosure Statement with Form PTO-1449
 Copies of IDS Citations
 Preliminary Amendment
 Return Receipt Postcard
 Small Entity Statement(s)
 Other:

Fee Calculation (37 CFR § 1.16)

(Col. 1) NO. FILED	(Col. 2) NO. EXTRA	SMALL ENTITY RATE FEE	OR	LARGE ENTITY RATE FEE
BASIC FEE		\$345 \$	OR	\$690 \$ 690
TOTAL CLAIMS	<u>27</u> -20 = <u>7</u>	x 9 = \$	OR	x 18 = \$ 126
INDEP CLAIMS	<u>10</u> -03 = <u>7</u>	x 39 = \$	OR	x 78 = \$ 546
[] Multiple Dependent Claim Presented		\$130 = \$	OR	\$260 = \$
* If the difference in Col. 1 is less than zero, enter "0" in Col. 2.		Total \$	OR	Total \$ 1,362

Check No. 2638 in the amount of \$ 1,402.00 is enclosed.

The Commissioner is authorized to charge any fees beyond the amount enclosed which may be required, or to credit any overpayment, to Deposit Account No. 50-0388 (Order No. CISCP141).

General Authorization for Petition for Extension of Time (37 CFR §1.136)

Applicants hereby make and generally authorize any Petitions for Extensions of Time as may be needed for any subsequent filings. The Commissioner is also authorized to charge any extension fees under 37 CFR §1.17 as may be needed to Deposit Account No. 50-0388 (Order No. CISCP141).

Please send correspondence to the following address:

Customer Number 022434
BEYER WEAVER THOMAS & NGUYEN, LLP
P.O. Box 130
Mountain View, CA 94042-0130
Telephone (510) 843-6200
Fax (510) 843-6203



Date: March 16, 2000

Joseph M. Villeneuve

Joseph M. Villeneuve
Registration No. 37,460

PATENT APPLICATION

METHODS AND APPARATUS
FOR REDIRECTING NETWORK TRAFFIC

Inventor(s): Shmuel Shaffer
1211 Cowper Street
Palo Alto, CA 94301
Citizenship: USA

James A. Aviani, Jr.
170 Olive Mill Road
Santa Barbara, CA 93108
Citizenship: USA

Alexander G. Tweedly
Maelstroma Kilmelford
Argyll PA34 4XD
Scotland
Citizenship: United Kingdom

Assignee: Cisco Technology, Inc.

BEYER & WEAVER, LLP
P. O. Box 130
Mountain View, California 94042-0130
(510) 843-6200

DRAFT - 08022005

PATENT

Attorney Docket No. CISCP141

5 METHODS AND APPARATUS
FOR REDIRECTING NETWORK TRAFFIC

BACKGROUND OF THE INVENTION

The present invention relates to transmission of data in a network environment.

More specifically, the present invention relates to methods and apparatus redirecting network traffic. Still more specifically, techniques are described herein for replicating packet flows 10 for a variety of purposes including, for example, troubleshooting, load balancing, and reliability.

Generally speaking, when a client platform communicates with some remote server, whether via the Internet or an intranet, it crafts a data packet which defines a TCP connection between the two hosts, i.e., the client platform and the destination server. More 15 specifically, the data packet has header fields which include the destination IP address, the destination port, the source IP address, the source port, and the protocol type. The destination IP address might be the address of a well known World Wide Web (WWW) search engine such as, for example, Yahoo, in which case, the protocol would be TCP and the destination port would be port 80, a well known port for http and the WWW. The source 20 IP address would, of course, be the IP address for the client platform and the source port would be one of the TCP ports selected by the client. These five pieces of information define the TCP connection.

Given the increase of traffic on the World Wide Web and the growing bandwidth demands of ever more sophisticated multimedia content, there has been constant pressure to 25 find more efficient ways to service data requests than opening direct TCP connections

between a requesting client and the primary repository for the desired data. Interestingly, one technique for increasing the efficiency with which data requests are serviced came about as the result of the development of network firewalls in response to security concerns. In the early development of such security measures, proxy servers were employed as firewalls to protect networks and their client machines from corruption by undesirable content and unauthorized access from the outside world. Proxy servers were originally based on Unix machines because that was the prevalent technology at the time. This model was generalized with the advent of SOCKS which was essentially a daemon on a Unix machine. Software on a client platform on the network protected by the firewall was specially configured to communicate with the resident demon which then made the connection to a destination platform at the client's request. The demon then passed information back and forth between the client and destination platforms acting as an intermediary or "proxy".

Not only did this model provide the desired protection for the client's network, it gave the entire network the IP address of the proxy server, therefore simplifying the problem of addressing of data packets to an increasing number of users. Moreover, because of the storage capability of the proxy server, information retrieved from remote servers could be stored rather than simply passed through to the requesting platform. This storage capability was quickly recognized as a means by which access to the World Wide Web could be accelerated. That is, by storing frequently requested data, subsequent requests for the same data could be serviced without having to retrieve the requested data from its original remote source. Currently, most Internet service providers (ISPs) accelerate access to their web sites using proxy servers.

Unfortunately, interaction with such proxy servers is not transparent, requiring each end user to select the appropriate proxy configuration in his or her browser to allow the browser to communicate with the proxy server. For the large ISPs with millions of

customers there is significant overhead associated with handling tech support calls from customers who have no idea what a proxy configuration is. Additional overhead is associated with the fact that different proxy configurations must be provided for different customer operating systems. The considerable economic expense represented by this overhead offsets the benefits derived from providing accelerated access to the World Wide Web. Another problem arises as the number of WWW users increases. That is, as the number of customers for each ISP increases, the number of proxy servers required to service the growing customer base also increases. This, in turn, presents the problem of allocating packet traffic among multiple proxy servers.

Network caching represents an improvement over the proxy server model which is transparent to end users, high performance, and fault tolerant. By altering the operating system code of an existing router, the router is enabled to recognize and redirect data traffic having particular characteristics such as, for example, a particular protocol intended for a specified port (e.g., TCP with port 80), to one or more network caches connected to the router via an interface having sufficient bandwidth. If there are multiple caches connected to the cache-enabled router, the router selects from among the available caches for a particular request based on a load balancing mechanism.

The network cache to which the request is re-routed “spoofs” the requested destination platform and accepts the request on its behalf via a standard TCP connection established by the cache-enabled router. If the requested information is already stored in the cache it is transmitted to the requesting platform with a header indicating its source as the destination platform. If the requested information is not in the cache, the cache opens a direct TCP connection with the destination platform, downloads the information, stores it for future use, and transmits it to the requesting platform. All of this is transparent to the user at the requesting platform which operates exactly as if it were communicating with the

destination platform. Thus, the need for configuring the requesting platform to suit a particular proxy configuration is eliminated along with the associated overhead. An example of such a network caching technique is embodied in the Web Content Caching Protocol (WCCP) provided by Cisco Systems, Inc., a specific embodiment of which is described in copending, commonly assigned, U.S. Patent Application No. 08/946,867 for METHOD AND APPARATUS FOR FACILITATING NETWORK DATA TRANSMISSIONS filed October 8, 1997, the entirety of which is incorporated herein by reference for all purposes.

Another specific embodiment of a packet redirection protocol which may be used to implement such a network caching technique is described in copending, commonly assigned,

- 10 U.S. Provisional Patent Application No. 60/168,862 for METHOD AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC filed December 2, 1999, the entirety of which is incorporated herein by reference for all purposes. According to a specific embodiment described in that application, the network caches have the capability of determining that particular redirected packets should be transmitted back to the redirecting router and
15 reinserted into the original traffic flow. This may be done in a manner transparent to the source or destination of the packets. An example of a case in which packets would need to be reinserted in the original flow might be where the cache recognizes the source and destination pairs identified by the packets as corresponding to a connection requiring IP-based authentication. Another example would be where the cache is overloaded and is
20 currently unable to handle all of the redirected traffic.

When information in a packet flow between two devices, e.g., a router and a network cache, is corrupted, it is useful for troubleshooting purposes to examine the sequence of packets as they are transmitted over the link. One technique for determining the cause of such corruption involves the insertion of a hub between the two devices and the connection
25 of a network sniffer to the hub. The hub duplicates all the packets in the flow and the sniffer

allows a technician to view the packet sequence. Unfortunately, this troubleshooting technique requires physically connecting the hub in close proximity to one of the two devices. As will be understood, this may not be practicable at geographically remote or isolated customer sites. Therefore, a technique is needed by which the troubleshooting of
5 packet flows in particular routers may be made more practicable.

Another problem associated with network transmissions is reliability. That is, because of a variety of network conditions (e.g., hardware and software failures, network congestion, etc.), transmitted packets occasionally fail to reach their intended destinations. When a packet flow encounters such a condition, a new packet flow must typically be
10 established to avoid the condition and make the specified connection. Not only is there a noticeable latency involved with reestablishing the packet flow, but there are certain client/server connections for which such an interruption is completely unacceptable. Therefore, there is also a need for improving the reliability with which data traffic is transmitted in network environments.

SUMMARY OF THE INVENTION

According to the present invention, a packet redirection technique is provided which allows a device to logically connect with a router and to receive all or part of a packet flow from an arbitrary network distance. According to a specific embodiment, the device connects with the router via any one of a variety of protocols including, for example, WCCP versions 1 and 2, specific implementations of which are described in the above-referenced copending patent applications. The registration process to connect the device to the router involves contacting the network administrator for the network which includes the router and requesting reconfiguration of the router to provide access to the router from a specific IP address corresponding to the device.

However, instead of redirecting the packet traffic to the logically connected device, the router instead replicates the packets of interest and transmits the replicate packets to the device via the redirection protocol. According to a specific embodiment, the characteristic(s) by which the packets of interest are identified are communicated to the router by the logically connected device. That is, when the remote device requests the connection with the router, it also requests the packets of interest by identifying at least one criterion by which the traffic may be identified. It will be understood that packet replication is preferable to actual redirection of the packets which could introduce an unacceptable latency into the packet flow when reinserted. According to a more specific embodiment, the packet redirection protocol with which the device connects to the router operates on the inbound as well as the outbound interfaces of the router so that both directions of a particular packet flow may be received by the logically connected device.

The present invention allows a remotely located engineer to investigate a network problem by actually inspecting the packet flow without disturbing the packet flow (and possibly masking the problem). And, as mentioned above, regardless of the distance

2000-04-26 10:30:20

between the engineer's test setup and the router under investigation, no additional latency is introduced into the packet flow.

In addition, the technique described herein may be used to increase the reliability with which certain traffic is transmitted in a network. That is, there are situations where in
5 order to ensure a high level of reliability for a particular connection, it would be useful to replicate the corresponding packet flow on two or more routing paths simultaneously.

Therefore, the packet replication of the present invention is employed as described above to send one or more replicate packet flows to devices logically connected to a particular router using, for example, WCCP. This ensures that the packets will reach the specified destination
10 in a timely manner.

According to yet another embodiment, the present invention may be used to implement a wide area load balancing scheme in which a client request to one of a plurality of equivalent servers is replicated to some or all of the plurality of servers for servicing by the server which responds to the client first. This ensures the quickest response to client
15 requests by the server most able to handle the traffic. It will be understood that this and the other innovations described herein may exist in combination with other load balancing or load distribution schemes such as, for example, those already existing in WCCP.

Thus, the present invention provides methods and apparatus for replicating a plurality of original packets in a packet flow which follows a first routing path. The packet flow is
20 received with a first device, the first device being included in the first routing path. In the first device, the original packets in the packet flow are identified according to at least one predetermined criterion. In the first device, replicate packets corresponding to the original packets are generated. The original packets are transmitted from the first device along the first routing path. The replicate packets are transmitted from the first device along a second
25 routing path which is different from the first routing path.

A further understanding of the nature and advantages of the present invention may be realized by reference to the remaining portions of the specification and the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a network diagram illustrating a troubleshooting technique according to a specific embodiment of the present invention;

Fig. 2 is a flowchart illustrating a troubleshooting technique according to a specific embodiment of the present invention;

Fig. 3 is a network diagram illustrating a packet replication technique for improved transmission reliability according to a specific embodiment of the present invention;

Fig. 4 is a flowchart illustrating a packet replication technique for improved transmission reliability according to a specific embodiment of the present invention;

Fig. 5 is a network diagram illustrating a load balancing technique according to a specific embodiment of the present invention;

Fig. 6 is a flowchart illustrating a load balancing technique according to a specific embodiment of the present invention; and

Fig. 7 is a simplified block diagram of a router for use with the various embodiments of the present invention.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

Fig. 1 is a simplified network diagram which will be used in conjunction with the flowchart of Fig. 2 to describe a specific embodiment of the present invention. According to this embodiment, a client machine 102 which is resident on a local area network (LAN) 104 communicates via router 106 and wide area network (WAN) 108, e.g., the Internet, with server 110. As discussed above, when the packet stream between client 102 and server 110 is corrupted, what typically happens is the network administrator of LAN 104 requests the assistance of an engineer employed by the vendor of router 106 which connects LAN 104 to the Internet. That is, the engineer actually goes to the geographic location of router 106, physically attaches a packet sniffer or equivalent test apparatus, and observes the packet flow to determine the cause of the corruption. The reason for this is that for many kinds of network problems it is critical to see the actual packet traffic. As will be understood, depending upon the accessibility of the geographic location of LAN 104, this can be an extremely expensive and inefficient exercise.

By contrast, the present invention allows a test device 112, e.g., a packet sniffer, to logically connect with router 106 from an arbitrary distance via a packet redirection protocol and remotely monitor all or a specified subset of the traffic through router 106. According to a specific embodiment, the protocol with which device 112 connects with router 106 is the Web Cache Coordination Protocol (WCCP) from Cisco Systems, Inc. Specific embodiments of WCCP are described in the above-referenced copending patent applications. According to a more specific embodiment, the protocol is WCCP v.2, i.e., WCCP2. The party desiring access to a particular router requests access from the network administrator over the phone or any other appropriate communication means. The network administrator then configures the router to allow the connection.

COPIED 08/25/2000

Referring now to Fig. 2, in a specific embodiment, once the network administrator configures router 106 to accept an expected connection from the IP address of test device 112, router 106 receives a WCCP connect request from test device 112 (202). According to a specific embodiment, the connect request from test device 112 additionally requests to receive copies of packets destined to client 102 and server 110. Once the connection via WCCP is established (204), router 106 continues receiving its packet flow (206) and begins to identify the packets in the packet flow of interest to the operator of test device 112 (208).

That is, router 106 identifies the packets of interest according to one or more predetermined criteria such as, in this example, the fact that the packets correspond to communications between a particular source and destination pair. This may be determined, for example, with reference to the source and destination IP addresses in the packet header. Notwithstanding this specific example, it will be understood that any of a wide variety of criteria for identifying packets of interest may be employed without departing from the scope of the invention.

Once the packets of interest have been identified, they are replicated by router 106 (210). According to an alternative embodiment, the identification and/or the replication of the packets of interest may occur before establishment of the connection with the test device. That is, router 106 could buffer such packets for later transmission to the test device. Once the packets of interest have been replicated, they are transmitted by router 106 along their original routing path to either server 110 or client 102 depending upon their direction (212). In such an embodiment, the packet redirection protocol, e.g., WCCP, would be working on both the inbound and outbound interfaces of router 106. In addition, the packet redirection may occur in a manner which is completely transparent to either end of the client/server communication.

The replicate packets are transmitted by router 106 to test device 112 along a different routing path (214) using the packet redirection protocol. If test device 112 is geographically remote relative to router 106, there may be a significant latency associated with the redirection, i.e., the packets may be seen at test device 112 much later than they are received at client 102 or server 110. Nevertheless, the engineer at test device 112 will see the sequence of packets as they occurred without having to travel to the location of LAN 104. As will be understood, this is an extremely valuable capability even where LAN 104 is within a few miles of the engineer's location.

Fig. 3 is another simplified network diagram which will be used in conjunction with the flowchart of Fig. 4 to describe another specific embodiment of the present invention. According to this embodiment, a client machine 302 which is resident on a LAN 304 may communicate with server 310 via router 306 and either WAN 308 or WAN 309, each of which represents one or more distinct routing paths between client 302 and server 310. It will be understood that WANs 308 and 309 may together represent the Internet.

According to a specific embodiment, router 306 replicates specific packets in a packet flow (such as the packets transmitted between client 302 and server 310) and transmits the replicate packets along one or more routing paths which are different than the original packets (e.g., via WAN 309 vs. WAN 308) for the purpose of improving the reliability of the connection between client 302 and server 310. In the illustrated embodiment, a router 312 is shown as part of WAN 309 and is logically connected to router 306 using WCCP2. It will be understood that device 312 may be any of a variety of network devices and still remain within the scope of the present invention. For example, device 312 could be a network cache, a work station, a file server, a switch, etc. It will also be

understood that the goal of reliability would be served if, for example, device 312 was serviced by a different Internet service provider (ISP) than the devices of WAN 308.

As router 306 receives its packet flow (402) it identifies packets of interest (404) as defined by the original request from router 312 by which the logical connection between the two routers was established. That is, the packets of interest in the packet flow are identified by router 306 in a manner similar to that described above with reference to Figs. 1 and 2.

The packets of interest are then replicated by router 306 (406). The original packets are transmitted to server 310 along the original routing path (408) while the replicate packets are transmitted to server 310 via at least one other alternate routing path (410). That is, the replicate packets are transmitted to router 312 which, in turn, transmits them to server 310. This replication and redirection may occur any number of times to provide additional levels of redundancy. Alternatively, the packets may be replicated more than once in router 306 and sent to any number of attached devices for forwarding to server 312.

It will be apparent that because of the redundancy provided by the redirected replicated traffic, the chances that the communication between client 302 and server 310 will be successful are greatly increased.

Fig. 5 is another simplified network diagram which will be used in conjunction with the flowchart of Fig. 6 to describe yet another specific embodiment of the present invention. According to this embodiment, a client machine 502 which is resident on a LAN 504 may communicate with servers 510-512 via router 506 and WAN 508, and servers 513 and 514 via router 506 and WAN 509. It will be understood that WANs 508 and 509 may together represent the Internet. Alternatively, WANs 508 and 509 may include any number of LAN/WAN configurations.

According to the embodiment of Figs. 5 and 6, a load balancing scheme is provided by which the packets from client 502 are replicated by router 506 and transmitted to a

plurality of equivalent servers 510-514. According to a specific embodiment, the first one of servers 510-514 to respond to a request from client 502 “wins” the connection. According to such an embodiment, all subsequent related traffic would then be transmitted between client 502 and the winning server. This would tend to favor creating connections with the server having the lightest load and/or using the routing path having the least congestion. That is, a wide area load balancing scheme is provided which has the ability to adjust to server and network traffic loads on the fly. According to a specific embodiment, each of servers 510-514 is logically connected to router 506 using a packet redirection protocol such as, for example, WCCP2. According to a more specific embodiment, in establishing its logical connection with router 506, each of servers 510-514 requests copies of packets directed to any of servers 510-514.

In the course of receiving and its normal packet flow (602), router 506 identifies the packets of interest in the packet flow (604) according to the criterion or criteria specified by the logically connected devices, i.e., servers 510-514. As discussed above, the packets of interest may be identified using a wide variety of criteria such as, in this example, the destination address(es) to which the packets are directed. That is, according to a specific embodiment, any packets addressed to any of servers 510-514 are identified in the packet flow. Once the packets of interest in the packet flow are so identified, router 506 replicates these packets (606) for transmission to the unspecified servers. The original packets of interest are then transmitted to the originally specified one of servers 510-514 along the original routing path (608), while the replicate packets are transmitted to the unspecified ones of servers 510-514 (610). As discussed above and according to one embodiment, the first one of servers 510-514 to respond to client 502 “wins” the connection. In this way, requests from client 502 will be serviced by the most accessible and available one of servers 510-514.

Generally, the packet redirection techniques of the present invention may be implemented on software and/or hardware. For example, each of the described techniques can be implemented in an operating system kernel, in a separate user process, in a library package bound into network applications, on a specially constructed machine, or on a network interface card. According to specific embodiments, the techniques of the present invention are implemented in software such as an operating system or in an application running on an operating system.

Software or software/hardware hybrid implementations of the invention may be implemented on general-purpose programmable machines selectively activated or reconfigured by a computer program(s) stored in memory. Such programmable machines may be a network device designed to handle network traffic. Such network devices typically have multiple network interfaces including frame relay and ISDN interfaces, for example. Specific examples of such network devices include routers and switches. For example, embodiments of the present invention may be implemented on specially configured routers such as specially configured router models 1600, 2500, 2600, 3600, 4500, 4700, 7200, 7500, and 12000 available from Cisco Systems, Inc. of San Jose, California. A general architecture for some of these machines will appear from the description given below. In an alternative embodiment, the present invention may be implemented on a general-purpose network host machine such as a personal computer or workstation. Further, the invention may be at least partially implemented on a card (e.g., an interface card) for a network device or a general-purpose computing device.

Referring now to Figure 7, a router 710 suitable for implementing the present invention includes a master central processing unit (CPU) 762, interfaces 768, and a bus 715 (e.g., a PCI bus). When acting under the control of appropriate software or firmware, the CPU 762 is responsible for such router tasks as routing table computations and network

management. It may also be responsible for facilitating connection to other device using a packet redirection protocol (e.g., WCCP2), identifying packets of interest, replication of packets, and transmission of packets to their respective destinations, etc. It preferably accomplishes all these functions under the control of software including an operating system

- 5 (e.g., the Internetwork Operating System (IOS®) of Cisco Systems, Inc.) and any appropriate applications software. CPU 762 may include one or more processors 763 such as a processor from the Motorola family of microprocessors or the MIPS family of microprocessors. In an alternative embodiment, processor 763 is specially designed hardware for controlling the operations of router 710. In a specific embodiment, a memory
- 10 761 (such as non-volatile RAM and/or ROM) also forms part of CPU 762. However, there are many different ways in which memory could be coupled to the system. Memory block 761 may be used for a variety of purposes such as, for example, caching and/or storing data, programming instructions, etc.

The interfaces 768 are typically provided as interface cards (sometimes referred to as "line cards"). Generally, they control the sending and receiving of data packets over the network and sometimes support other peripherals used with the router 710. Among the interfaces that may be provided are Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like. In addition, various very high-speed interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM. The independent processors may control such communications intensive tasks as packet switching, media control and management. By providing separate processors for the communications intensive tasks, these interfaces allow

the master microprocessor 762 to efficiently perform routing computations, network diagnostics, security functions, etc.

Although the system shown in Figure 7 is one specific router of the present invention, it is by no means the only router architecture on which the present invention can be implemented. For example, an architecture having a single processor that handles communications as well as routing computations, etc. is often used. Further, other types of interfaces and media could also be used with the router.

Regardless of network device's configuration, it may employ one or more memories or memory modules (such as, for example, memory block 765) configured to store data, program instructions for the general-purpose network operations and/or the packet redirection and replication functions described herein. The program instructions may control the operation of an operating system and/or one or more applications, for example. The memory or memories may also be configured to store packets for replication, replicated packets, packet identification criteria, etc.

Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). The invention may also be embodied in a carrier wave travelling over an appropriate medium such as airwaves, optical lines, electric lines, etc. Examples of program instructions include both machine

code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

While the invention has been particularly shown and described with reference to specific embodiments thereof, it will be understood by those skilled in the art that changes in the form and details of the disclosed embodiments may be made without departing from the spirit or scope of the invention. For example, specific embodiments of the invention have been described herein with reference to the use of WCCP for logically connecting a router to another device to which replicate packets are then transmitted. It will be understood, however, that other protocols may be employed to establish the alternate routing path by which replicate packets are transmitted without departing from the scope of the invention.

Moreover, the techniques of the present invention have been described in the context of troubleshooting, reliability, and load balancing schemes. It will be understood, however, that a wide variety of application could advantageously employ the techniques described herein. For example, transparent packet monitoring could be implemented from remote locations. In addition, the specific embodiments described herein have been described with reference to particular network topologies. However, it will be understood that the techniques described herein may be implemented on any of a wide variety of network topologies without departing from the scope of the invention. Therefore, the scope of the invention should be determined with reference to the appended claims.

WHAT IS CLAIMED IS:

1. A method for replicating a plurality of original packets in a packet flow, the packet flow following a first routing path, the method comprising:

5 receiving the packet flow with a first device, the first device being included in the first routing path;

in the first device, identifying the original packets in the packet flow according to at least one predetermined criterion;

10 in the first device, generating replicate packets corresponding to the original packets;

transmitting the original packets from the first device along the first routing path; and

transmitting the replicate packets from the first device along a second routing path, the second routing path being different from the first routing path.

2. The method of claim 1 wherein the second routing path includes a second

15 device, the second device being logically connected with the first device via a protocol.

3. The method of claim 2 wherein the protocol comprises a packet redirection protocol.

20 4. The method of claim 3 wherein the packet redirection protocol comprises an object caching protocol.

5. The method of claim 2 wherein the original packets indicate a destination device, the destination device being included in the first routing path, the first device

25 transmitting the original packets to the destination device via the first routing path, the

second device facilitating transmission of the replicate packets to the destination device via the second routing path.

6. The method of claim 2 wherein the second device comprises a test device for facilitating inspection of the replicate packets.

7. The method of claim 1 wherein each of the original packets indicate one of a plurality of destination devices each of the destination devices being logically connected with the first device via a protocol, a first one of the destination devices being included in the first routing path, a second one of destination devices being included in the second routing path, and wherein the replicate packets are transmitted along the second routing path to the second one of the destination devices.

8. The method of claim 7 further comprising:
 - determining which of the original and replicate packets reach their respective destination devices first, thereby identifying a winner destination device; and
 - awarding a connection to an originating device to the winner destination device.

9. The method of claim 2 further comprising:
 - receiving a request from the second device for connecting with the first device via the protocol; and
 - connecting with the second device via the protocol.

10. The method of claim 9 wherein the request from the second device identifies the at least one predetermined criterion.

11. The method of claim 1 wherein the original packets originate from a source device, the method for replicating the original packets being transparent to the source device.

5 12. The method of claim 1 wherein the original packets indicate a destination device, the method for replicating the original packets being transparent to the destination device.

10 13. The method of claim 1 wherein the first device comprises a router.

14. The method of claim 1 wherein the at least one predetermined criterion comprises a source address.

15 15. The method of claim 1 wherein the at least one predetermined criterion comprises a destination address.

20 16. The method of claim 1 wherein the at least one predetermined criterion comprises a socket.

17. The method of claim 1 wherein the at least one predetermined criterion comprises a port.

25 18. The method of claim 1 wherein the at least one predetermined criterion comprises a protocol type.

19. An apparatus for replicating a plurality of original packets in a packet flow, the packet flow following a first routing path, the method comprising:

means for receiving the packet flow, the receiving means being included in the first routing path;

5 means for identifying the original packets in the packet flow according to at least one predetermined criterion;

means for generating replicate packets corresponding to the original packets;

means for transmitting the original packets along the first routing path; and

means for transmitting the replicate packets along a second routing path, the second

10 routing path being different from the first routing path.

20. A router operable to replicate a plurality of original packets in a packet flow, the packet flow following a first routing path, the router comprising:

a memory having at least a portion of a router operating system stored therein; and

15 a processor for controlling operation of the router according to the router operating system, the processor being configured by the router operating system to:

receive the packet flow with the router, the router being included in the first routing path;

identify the original packets in the packet flow according to at least

20 one predetermined criterion;

generate replicate packets corresponding to the original packets;

transmit the original packets from the router along the first routing path; and

transmit the replicate packets from the router along a second routing path, the second routing path being different from the first routing path.

21. A computer program product for replicating a plurality of original packets in a packet flow, the packet flow following a first routing path, the computer program product comprising:

5 at least one computer readable medium; and
computer program instructions stored in the at least one computer readable medium
for causing a processing device to:

receive the packet flow, the processing device being included in the
first routing path;

10 identify the original packets in the packet flow according to at least
one predetermined criterion;

generate replicate packets corresponding to the original packets;

transmit the original packets along the first routing path; and

transmit the replicate packets along a second routing path, the second
routing path being different from the first routing path.

15
22. A method for remotely monitoring a portion of a packet flow associated with
a first device using a second device, the packet flow following a first routing path, the
method comprising:

20 receiving a request from the second device for connecting with the first device via a
protocol;

logically connecting with the second device via the protocol;

receiving the packet flow with the first device, the first device being included in the
first routing path;

in the first device, identifying original packets in the packet flow according to at least one predetermined criterion;

in the first device, generating replicate packets corresponding to the original packets;

transmitting the original packets from the first device along the first routing path; and

transmitting the replicate packets from the first device to the second device along a

second routing path, the second routing path being different from the first routing path.

23. A router operable to facilitate monitoring by a remote device of a portion of a packet flow associated with a router, the packet flow following a first routing path which

includes the router, the router comprising:

a memory having at least a portion of a router operating system stored therein; and

a processor for controlling operation of the router according to the router operating system, the processor being configured by the router operating system to:

receive a request from the remote device for connecting with the first device via a protocol;

logically connect with the remote device via the protocol;

receive the packet flow;

identify original packets in the packet flow according to at least one predetermined criterion;

20 generate replicate packets corresponding to the original packets;

transmit the original packets along the first routing path; and

transmit the replicate packets to the remote device along a second routing path, the second routing path being different from the first routing path.

24. A method for replicating a plurality of original packets in a packet flow, the packet flow following a first routing path, the original packets indicating a destination device, an intermediate device and the destination device being included in the first routing path, the method comprising:

5 receiving the packet flow with the intermediate device;

in the intermediate device, identifying the original packets in the packet flow according to at least one predetermined criterion;

in the intermediate device, generating replicate packets corresponding to the original packets;

10 transmitting the original packets from the intermediate device to the destination device along the first routing path; and

transmitting the replicate packets from the intermediate device to the destination device along at least one other routing path, the at least one other routing path being different from the first routing path.

15

25. A router operable to replicate a plurality of original packets in a packet flow, the packet flow following a first routing path, the original packets indicating a destination device, the router and the destination device being included in the first routing path, the router comprising:

20 a memory having at least a portion of a router operating system stored therein; and
a processor for controlling operation of the router according to the router operating system, the processor being configured by the router operating system to:

receive the packet flow;

identify the original packets in the packet flow according to at least

25 one predetermined criterion;

- generate replicate packets corresponding to the original packets;
- transmit the original packets to the destination device along the first routing path; and
- transmit the replicate packets to the destination device along at least one other routing path, the at least one other routing path being different from the first routing path.

- 5 26. A method for replicating a plurality of original packets in a packet flow, the packet flow following a first routing path, the original packets indicating one of a plurality of 10 destination devices, an intermediate device and a first one of the destination devices being included in the first routing path, the method comprising:
- receiving the packet flow with the intermediate device;
- in the intermediate device, identifying the original packets in the packet flow according to at least one predetermined criterion;
- 15 in the intermediate device, generating replicate packets corresponding to the original packets;
- transmitting the original packets from the intermediate device to the first destination device along the first routing path; and
- transmitting the replicate packets from the intermediate device to at least one other of 20 the destination devices along at least one other routing path.

- 25 27. A router operable to replicate a plurality of original packets in a packet flow, the packet flow following a first routing path, the original packets indicating one of a plurality of destination devices, the router and a first one of the destination devices being included in the first routing path, the router comprising:

DOCUMENT - 58004560

a memory having at least a portion of a router operating system stored therein; and
a processor for controlling operation of the router according to the router operating
system, the processor being configured by the router operating system to:

receive the packet flow;

5 identify the original packets in the packet flow according to at least
one predetermined criterion;

generate replicate packets corresponding to the original packets;

transmit the original packets to the first destination device along the
first routing path; and

10 transmit the replicate packets to at least one other of the destination
devices along at least one other routing path.

00000000000000000000000000000000

METHODS AND APPARATUS
FOR REDIRECTING NETWORK TRAFFIC

ABSTRACT OF THE DISCLOSURE

5 Methods and apparatus are described for replicating a plurality of original packets in a packet flow which follows a first routing path. The packet flow is received with a first device, the first device being included in the first routing path. In the first device, the original packets in the packet flow are identified according to at least one predetermined criterion. In the first device, replicate packets corresponding to the original packets are 10 generated. The original packets are transmitted from the first device along the first routing path. The replicate packets are transmitted from the first device along a second routing path which is different from the first routing path.

DRAFTED - 2005/02/25 00:00:00

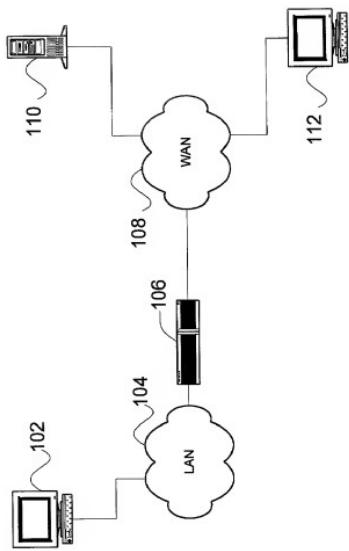


Fig. 1

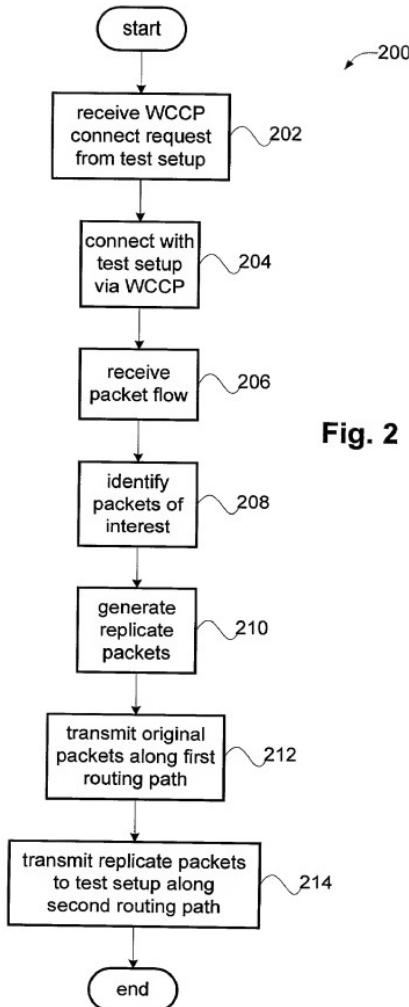


Fig. 2

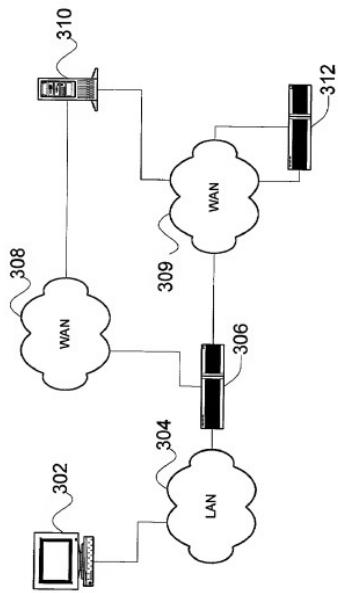


Fig. 3

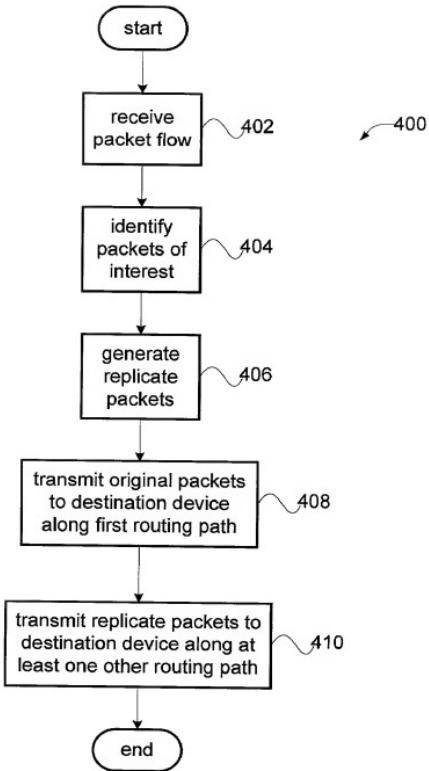


Fig. 4

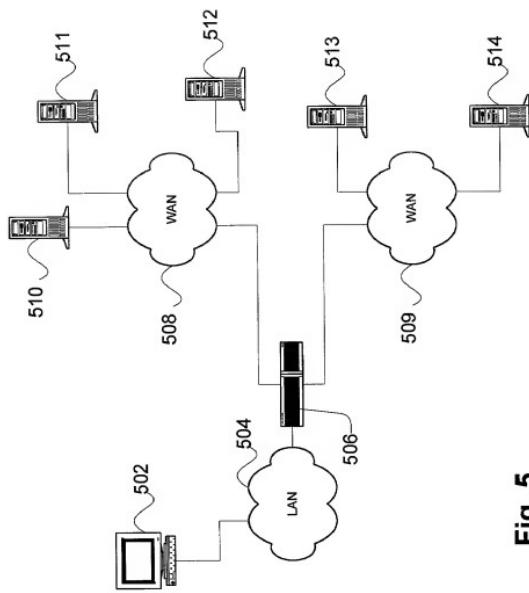


Fig. 5

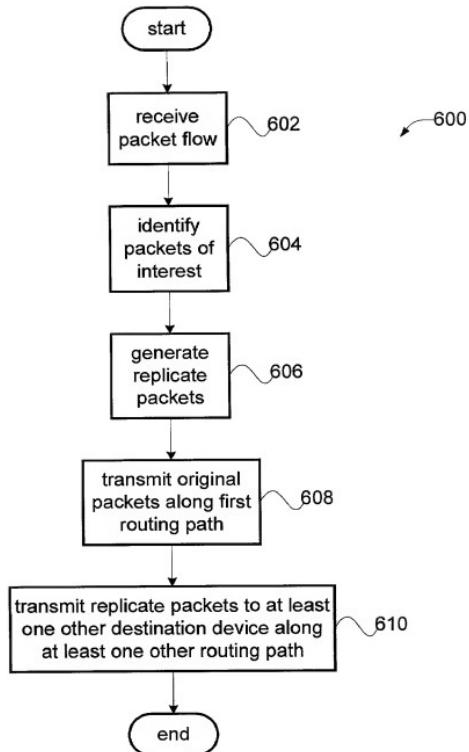


Fig. 6

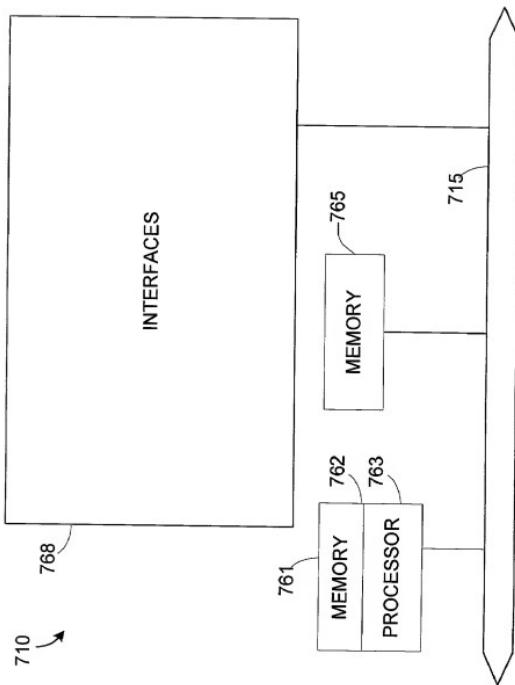


Fig. 7

**DECLARATION AND POWER OF ATTORNEY
FOR ORIGINAL U.S. PATENT APPLICATION**

Attorney's Docket No. CISCP141

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **METHODS AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC**, the specification of which,

(check one) 1. is attached hereto.

2. was filed on _____ as
U.S. Application No. _____
and was amended on _____.

3. was filed on _____ as
International PCT Application No. _____
and was amended on _____.

hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, CFR § 1.56.

hereby claim foreign priority benefits under Title 35, United States code, § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Priority Benefits Claimed?

(Appl. No.)

(Country)

(Filing Date)

Yes No

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below:

Prior Provisional Application(s)

(Application No.)

(Filing Date)

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s), or § 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

Prior U.S. Application(s)

(Application No.)

(Filing Date)

(Status - patented, pending, abandoned)

And I hereby appoint the law firm of **Beyer & Weaver, LLP** and all practitioners who are associated with the Customer Number 022434 as my principal attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

Direct Correspondence To:

Customer Number: 022434
BEYER & WEAVER, LLP
P.O. Box 130
Mountain View, CA 94042-0130



Direct Telephone Calls To:

Attorney at telephone number (510) 843-6200

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

Typewritten Full Name of
First Joint Inventor:

Shmuel Shaffer

Citizenship: USA

Inventor's signature:

Date of Signature: 3/15/16

Residence: (City)

Palo Alto

(State/Country) CA/USA

Post Office Address:

1211 Cowper Street, Palo Alto, CA 94301

Typewritten Full Name of
Second Joint Inventor:

James A. Aviani, Jr.

Citizenship: USA

Inventor's signature:

Date of Signature:

Residence: (City)

Santa Barbara

(State/Country) CA/USA

Post Office Address:

170 Olive Mill Road, Santa Barbara, CA 93108

Typewritten Full Name of
Third Joint Inventor: Alexander G. Tweedly Citizenship: United Kingdom
Inventor's signature: _____ Date of Signature: _____
Residence: (City) Argyll (State/Country) Scotland
Post Office Address: Maelstroma Kilmelford, Argyll PA34 4XD, Scotland

**DECLARATION AND POWER OF ATTORNEY
FOR ORIGINAL U.S. PATENT APPLICATION**

Attorney's Docket No. **CISCP141**

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **METHODS AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC**, the specification of which,

(check one) is attached hereto.

2. was filed on _____ as
U.S. Application No. _____
and was amended on _____.

3. was filed on _____ as
International PCT Application No. _____
and was amended on _____.

hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, CFR § 1.56.

hereby claim foreign priority benefits under Title 35, United States code, § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Priority Benefits Claimed?

(Appl. No.)

(Country)

(Filing Date)

Yes No

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below:

Prior Provisional Application(s)

(Application No.)

(Filing Date)

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s), or § 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

Prior U.S. Application(s)

(Application No.)

(Filing Date)

(Status - patented, pending, abandoned)

And I hereby appoint the law firm of **Beyer & Weaver, LLP** and all practitioners who are associated with the Customer Number 022434 as my principal attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

Direct Correspondence To:

Customer Number: 022434
BEYER & WEAVER, LLP
P.O. Box 130
Mountain View, CA 94042-0130



Direct Telephone Calls To:

Attorney at telephone number (510) 843-6200

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

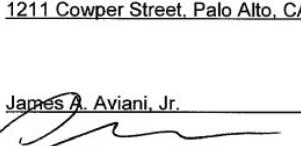
Typewritten Full Name of

First Joint Inventor:

Shmuel Shaffer

Citizenship: USA

Inventor's signature:



Date of Signature: _____

Residence: (City)

Palo Alto

(State/Country) CA/USA

Post Office Address:

1211 Cowper Street, Palo Alto, CA 94301

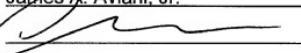
Typewritten Full Name of

Second Joint Inventor:

James A. Aviani, Jr.

Citizenship: USA

Inventor's signature:



Date of Signature: 3/13/2000

Residence: (City)

Santa Barbara

(State/Country) CA/USA

Post Office Address:

170 Olive Mill Road, Santa Barbara, CA 93108

Typewritten Full Name of
Third Joint Inventor: Alexander G. Tweedly Citizenship: United Kingdom
Inventor's signature: _____ Date of Signature: _____
Residence: (City) Argyll (State/Country) Scotland
Post Office Address: Maelstroma Kilmelford, Argyll PA34 4XD, Scotland

DUKE = 58072560

**DECLARATION AND POWER OF ATTORNEY
FOR ORIGINAL U.S. PATENT APPLICATION**

Attorney's Docket No. **CISCP141**

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **METHODS AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC**, the specification of which,

(check one)

1. is attached hereto.
2. was filed on _____ as U.S. Application No. _____ and was amended on _____.
3. was filed on _____ as International PCT Application No. _____ and was amended on _____.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, CFR § 1.56.

I hereby claim foreign priority benefits under Title 35, United States code, § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Priority Benefits Claimed?

(Appl. No.) (Country) (Filing Date)

Yes No

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below:

Prior Provisional Application(s)

(Application No.) (Filing Date)

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s), or § 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

Prior U.S. Application(s)

(Application No.) _____ (Filing Date) _____ (Status - patented, pending, abandoned)

(Filing Date)

(Status - patented, pending, abandoned)

And I hereby appoint the law firm of **Beyer & Weaver, LLP** and all practitioners who are associated with the Customer Number 022434 as my principal attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

Direct Correspondence To:

Customer Number: 022434
BEYER & WEAVER, LLP
P.O. Box 130
Mountain View, CA 94042-0130



Direct Telephone Calls To:

Attorney at telephone number (510) 843-6200

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

Typewritten Full Name of
First Joint Inventor:

Shmuel Shaffer

Citizenship: USA

Inventor's signature:

Date of Signature:

Residence: (City)

Palo Alto

(State/Country) CA/USA

Post Office Address:

1211 Cowper Street, Palo Alto, CA 94301

Typewritten Full Name of
Second Joint Inventor:

James A. Aviani, Jr.

Citizenship: USA

Inventor's signature:

Date of Signature:

Residence: (City)

Santa Barbara

(State/Country) CA/USA

Post Office Address:

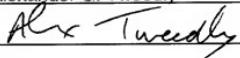
170 Olive Mill Road, Santa Barbara, CA 93108

Typewritten Full Name of
Third Joint Inventor:

Alexander G. Tweedly

Citizenship: United Kingdom

Inventor's signature:



Date of Signature: March 7, 2000

Residence: (City)

Argyll

(State/Country) Scotland

Post Office Address:

Maelstroma Kilmelford, Argyll PA34 4XD, Scotland